

Automatic Fetal Face Detection From Ultrasound Volumes Via Learning 3D and 2D Information

Shaolei Feng¹, S. Kevin Zhou¹, Sara Good², and Dorin Comaniciu¹

¹Integrated Data Systems Department, Siemens Corporate Research, Princeton, NJ 08540

²Siemens Medical Solutions, Innovations Division, CA 94043

Abstract

3D ultrasound imaging has been increasingly used in clinics for fetal examination. However, manually searching for the optimal view of the fetal face in 3D ultrasound volumes is cumbersome and time-consuming even for expert physicians and sonographers. In this paper we propose a learning-based approach which combines both 3D and 2D information for automatic and fast fetal face detection from 3D ultrasound volumes. Our approach applies a new technique – constrained marginal space learning – for 3D face mesh detection, and combines a boosting-based 2D profile detection to refine 3D face pose. To enhance the rendering of the fetal face, an automatic carving algorithm is proposed to remove all obstructions in front of the face based on the detected face mesh. Experiments are performed on a challenging 3D ultrasound data set containing 1010 fetal volumes. The results show that our system not only achieves excellent detection accuracy but also runs very fast – it can detect the fetal face from the 3D data in 1 second on a dual-core 2.0 GHz computer.

1. Introduction

Nowadays, ultrasound has been an important medical imaging modality for visualizing and diagnosing internal organs and fetus's. Compared to other modalities such as magnetic resonance imaging (MRI) and computed tomography (CT), ultrasound technology is safe, inexpensive and portable. The 3D ultrasound imaging is an extension of the standard 2D ultrasound imaging. The reflected echoes are processed by computer programs to reconstruct 3D volumetric images of the internal organs or fetus. Although 3D ultrasound is increasingly used in clinics for fetal examination, it is not easy to rapidly and precisely navigate to the fetal face surface in order to render an optimal face view. Even for trained sophisticated doctors, manually locating a fetal face in 3D/4D ultrasound data is challenging and tedious. In general, it takes about 8 to 10 minutes to navigate




Figure 1. Examples of 3D views of fetal faces from our ultrasound volumes. Note the complex background and large variation in appearance.

to a pleasant view of the fetal face for an expert sonographer.

In this paper, we propose a learning based approach which combines both 2D and 3D information for automatic fetal face detection in 3D ultrasound volumes. Aided by automatic detection, the long learning curve needed for 3D scanning can be reduced as the system automatically locates the acquisition plane to obtain an optimal view. In addition, based on the automatic detection results, one can omit anything in front the face to achieve better views. Quantitative measurements and analysis are also possible after the 3D automatic detection. However, automatic fetal face detection from 3D ultrasound data is a very challenging problem. The variation of the ultrasound fetus data is very large. The appearances of fetal faces at different pregnant stages vary a lot, even for the same fetus. The positions and poses of the fetuses in the scanned ultrasound data are also very different. The ultrasound volumes may have low image qualities, imaging artifacts like small particles, and blur or




Figure 2. Examples of the initial position of a loaded volume.

dark parts because of weak signals. Cord, placenta, uterus, and extremities can occlude the fetal face; Manual scanning may result in incomplete faces with the forehead or cheek missing from the volume. Besides these challenges arising from data acquisition, it is also very hard to acquire precise ground truth of the fetal faces for our supervised learning approach. For example, the peripheral outline of the face is hard to determine when annotating the data because of occlusions, missing face parts or the quality of multi-planar reformatting (MPR) planes, which introduces uncertainty.

Figure 1 shows some examples of the fetal face 3D views randomly selected from our data set, from which we can see the complex background, noise and blurs in the ultrasound volumes. Note that all these 3D views are rendered at the optimal position through manual navigation in the 3D volumes. The original loading positions of these volumes were not acquired at the optimal positions. Figure 2 shows an example of the original loading position of one volume, from which it is almost impossible to see the face.

To make the detection robust and precise, we propose combining both 3D face surface detection and 2D face profile detection to obtain an accurate 3D face detection. For the 3D face surface detection, we adopt a mesh based detection method which treats the detection problem as a two-stage task: rigid object localization and non-rigid boundary shape detection. Furthermore, when sonographers scan for the fetal face, they start the acquisition at the profile. So based on the observation that in 3D ultrasound data a 2D face profile is usually clearer and easier to acquire, we combine the 3D face surface detection with a 2D face profile detection which further refines the 3D detection. The combined method implicitly gives more weights to those

distinctive and stable features along the face profile and reduced the effect of noise and outliers from other parts of the face.

1.1. Related Work

Although a lot of work has been done in the 2D image face detection field (please refer to surveys [14, 15]), less work was done for 3D face detection [1, 3, 4, 6, 9, 10, 13]. Furthermore, to our best knowledge this is the first work to detect fetal faces from 3D ultrasound data. Most approaches for 2D face detection employ an exhaustive search over the input image, which is hard to directly extend to 3D face detection because the calculation complexity would exponentially increase with the dimensions of the parameter space. Colombo *et al* [4] proposed a curvature analysis based detection approach for 3D faces acquired by a laser range scanner. Their method first detects salient face features such as eyes and nose, and then through face surface curvature analysis a PCA-based classifier is applied to the detected candidate noses and eyes to determine if they are real faces. Wang *et al*. [13] proposed “point signature” representation for 3D surface and combined the 2D image Gabor features for face detection and recognition. Moreno *et al* [9] segmented faces using surface curvatures. Face surface analysis approaches based on face profile extraction are also studied by Cartoux *et al*. [3], Beumier and Acheroy [1], and Pan *et al*. [10].

This work is different with previous one in several aspects. First, our approach combines both 3D and 2D information to detect the fetal face. The 2D profile detection following the 3D surface detection makes the approach more robust. Second, we use learning based approaches for both the 3D face surface detection and the 2D face profile detection. Third, a new technique, the constrained marginal space learning proposed by [17] is used for 3D face surface detection. The 2D profile detection is conducted with a boosted selection of features [5, 11, 12]. Fourth, our system performs quite fast. It takes about one second to detect the 3D face on a dual-core 2.0 GHz computer.

2. Fetal Face Representation

We expect our automatic fetal face detection system to facilitate multiple purposes of medical fetal examination, e.g. measurement of particular anatomies and carving for better rendering views. This requires a relative precise representation of the fetal faces, i.e. besides the exact pose information of the fetal faces, the positions of important landmarks (e.g. eyes and nose) of the faces are also necessary. On the other hand, overly precise face representation introduces much work for manually annotating faces. In order to meet all the requirements, we proposed a relative sparse mesh representation $\mathcal{M}(P, T)$ for the fetal faces,




Figure 3. Mesh representation for a fetal face.

where $P = \{p_i \in \mathbb{R}^3\}_{i=1}^n$ is the set of n mesh points and $T = \{t_i \in \mathbb{Z}^3\}_{i=1}^m$ the set of m mesh triangles. In our fetal face mesh, the mesh point set P consists of three main components: face profile points, face contour points and other facial landmarks.

Face profile points are a set of key points along the face profile and face contour points define the peripheral outline for the face mesh. There are 8 other facial landmarks, including four eye points, two nose points and two mouth points. The points of each of the first two components – profile and contour – must be co-planar, and the two planes determined by the profile and the boundary respectively must be perpendicular to each other. These geometry constraints simplify the annotation procedure and also guarantee the consistence of annotations over large data sets. Figure 3 shows an example of an annotated face mesh for a fetal face volume, with the middle red line for the face profile, four magenta points for eyes, two cyan points for nose roots and two yellow points for mouth corners.

3. 3D Fetal Face Detection

The whole procedure of our fetal face detection follows the workflow shown in Figure 4, which consists of three steps: 3D face surface localization, face location refinement based on 2D profile detection and 3D face non-rigid deformation. In the 3D face localization step we find the bounding box of fetal face in 3D volumes, which is described by a set of rigid parameters – its position, orientation and scale. This step will generate multiple candidate locations of the fetal face. In the next step, a 2D face profile detector will be applied to refine the mean location aggregated from the set of candidates. Then a mean shape is aligned with the estimated location as a rough estimate of the face shape. In the final step, a 3D boundary detector is used to move each landmark to the optimal position and a non-rigid deformation of the mean shape is made towards the optimal face shape.

We employ a learning-based method for the 3D localization of fetal face, i.e. finding a bounding box for the fetal face. The learning-based method takes box detection as a two-category classification problem: each candidate box either contains the target object or not. The location of each candidate box is parameterized by 9 transformation parameters (3 for translation, 3 for orientation and 3 for scale). Although exhaustive search is widely used for 2D object detection for its robustness, for 3D objects it is extremely expensive to exhaustively consider all possible candidates. Under a uniform sampling scheme over the whole parameter space, the number of tested candidates exponentially increases with the dimensionality of the parameter space. In our fetal face detection system, we used a marginal space learning framework proposed by Zheng, *et al* [16], which basically learns the location parameters in a sequential way on projected sampling sub-spaces. A classifier is trained for each projected space using the probabilistic boosting-tree (PBT) [11].

For 2D face profile detection, we adopt the boosting models similar to those used in [5]. The models first detect the rigid transformation of the profile structure through an exhaustive search over the whole cutting image, then perform a shape inference to find the most likely associated profile shape.

In the following sections we will discuss the models we used for 3D face surface detection and 2D face profile detection.

3.1. Constrained Marginal Space Learning for 3D Fetal Face Surface Localization

In the 3D fetal face surface localization task, given a 3D ultrasound volume V we want the system to output a set of volume blocks each of which has a high probability of containing the fetal face. Denote the block parameter as $[p, \sigma, s] \in \mathbb{R}^9$, where $p = (X, Y, Z)$ represents the position, $\sigma = (\psi, \phi, \theta)$ the orientation and $s = (S_x, S_y, S_z)$ the scale along x, y, z axes respectively. Then this task is to find the set of blocks each of which satisfies:

$$Pr(p, \sigma, s|V) > thr, \quad (1)$$

where thr is the threshold determining if a block contains the target object. It is empirically selected by tuning on the training set. Marginal space learning (MSL) provides an efficient inference scheme to solve (1), which breaks the original parameter space Ω into subsets of increasing marginal space $\Omega_1 \subset \Omega_2 \subset \dots \subset \Omega_n = \Omega$ and sequentially train and detect over each subset. In fetal face detection, we assume the whole parameter space is decomposed into the following sequence: $\Omega_1 = p$, $\Omega_2 = [p, \sigma]$ and $\Omega_3 = [p, \sigma, s]$. Accordingly, we have:

$$Pr(p, \sigma, s|V) = Pr(p|V)Pr(\sigma|p, V)Pr(s|\sigma, p, V). \quad (2)$$




Figure 4. Diagram for 3D fetal face detection.

Each of the marginal probabilities is estimated by the probabilistic boosting tree (PBT). For position probability estimation, 3D Haar features are used and for orientation and scale probabilities steerable features [16] are used. In the marginal space learning, each marginal space is uniformly sampled with an assumption that the parameters of that marginal space are independent with each other. However, in many applications the parameters of a marginal space are not independent. Usually, there are correlations among their distributions. For example, in fetal face detection, the parameters in the orientation space $\sigma = (\psi, \phi, \theta)$ represent the three Euler angles of a fetal face rotated along x , y and z axes respectively. Although each parameter has a large variation range, the combined parameters σ are not uniformly distributed in the 3D orientation space. Figure 6.(c) shows the distribution of the σ in the 3D space, which indicates a strong correlation among the three Euler angles. Uniformly sampling this marginal space will include much more hypotheses than necessary, which slows down the training and detection process. Moreover, potential outliers introduced by uniformly sampling may hurt the detection performance. Similar situation also holds for the scale space.

To reduce the unnecessary hypotheses tested, we use a constrained marginal space learning framework proposed by [17], which proposes an example-based strategy to generate testing hypotheses. The basic procedure is simple: first uniformly sample the marginal space to get the initial hypothesis set H_u , and then for each training sample, search the neighboring hypotheses in the H_u and insert those hypotheses in the selected set H , finally remove all redundant hypotheses from H . Then H is the generated testing hypotheses for detection. Using the constrained marginal space learning, both the detection speed and accuracy are improved.

Another improvement of the constrained marginal space learning is that the Euler angle representation of 3D orientation $\sigma = (\psi, \phi, \theta)$ is converted to the unit quaternion representation $q = [s, \mathbf{w}]$ proposed by Karney et al. [7], where s is a scalar and \mathbf{w} a vector. Each quaternion is also expressed as

$$q = [\cos(\theta/2), \mathbf{v} \cdot \sin(\theta/2)], \quad (3)$$

where $\mathbf{v} \in \mathbb{R}$ and $|\mathbf{v}| = 1$. This basically means a rotation of θ around the axis \mathbf{v} . The original Euler angle representation for 3D orientation has two main drawbacks for

estimating orientation marginal probability. First, the Euler angle representation for a 3D orientation is not unique. Second, uniform step size of Euler angles does not generate a uniform sampling in the orientation space, which causes uncertainty of the distribution of the samples in the orientation space. The quaternion representation overcomes these drawbacks. One important property is that the composition of two rotations can be computed as the product of two quaternions, which makes the calculation of rotation distance from one orientation O_1 to another O_2 easier:

$$D(O_1, O_2) = \arccos(|s_1 s_2 - \mathbf{w}_1 \cdot \mathbf{w}_2|). \quad (4)$$

One can refer [7] for the details and [2] for an example of its applications in fetal head anatomy indexing. The quaternion representation guarantees that the initial sampling over the orientation space is uniform and avoids the orientation ambiguity of Euler angle representation.

After the candidates of the bounding box of the fetal face are detected, the mean shape of the fetal face mesh is aligned within each box as a rough shape of the candidate face.

3.2. 2D Face Profile Detection and Results Refinement

In our system, the 3D box detector generates a set of fetal face candidates. In this section we discuss the 2D face profile detection and how it is used to refine the detected 3D locations.

The intuition to refine the 3D mesh detection using a 2D profile detection is based on the following facts. First, in 3D ultrasound volumes of fetus the 2D face profile is usually clearer and easier to acquire comparing with the whole face surface. When sonographers acquire ultrasound volumes for fetal head, they usually first navigate the transducer to capture the face profile, serving as an important clue. Second, the profile detector is more robust and performs better than the 3D detection in terms of accuracy. The 2D profile detection is less sensitive to occlusions and missing face parts. For example, even if the whole cheek is missed or a hand covers the whole cheek of the face in the 3D volume, the profile will not be effected.

We adopt a simple way to refine the 3D mesh detection using the 2D profile detection, which does show the improvement on final face detection as experiments indicate:




Figure 5. Illustration of the detected 3D mesh from the MSL and the detected 2D profile for refinement. The detected mesh in the left picture is off the true position indicated by the detected 2D profile. Through refinement, it is dragged back to the correct position as the right picture shows.

1. Based on the pool of candidate boxes of the 3D box detection, we aggregate one mean box and associate it with a mean face mesh.
2. From this mean face mesh at the aggregated position, the initial profile points can be easily extracted and used to determine the MPR plane.
3. A 2D face profile detector is then applied to this MPR plane to detect a new profile.
4. Adjust the location and pose of the detected mean mesh according to the newly detected profile.

Figure 5 shows an example of the detected 3D mesh from the MSL and the detected 2D profile for refinement. The 3D mesh is off the true position, but through the 2D profile refinement it can be dragged back to the true position.

We also tried other approaches to refine the 3D mesh detection using the 2D profile detection. For example, instead of refining the aggregated mesh we refine each candidate from the MSL detection and select the one with the highest 2D detection probability. But this does not help the final detection in our experiments because the best profile image does not always cut through the exact middle of the face.

For 2D face profile detection, given an MPR I extracted from the volume the system needs to determine the face profile position and shape. This step is performed by a probabilistic-boosting-tree (PBT) based object localization and a nearest-neighboring based shape inference for profile shape optimization. Our PBT for profile localization is similar to the boosted cascade of Haar features by Viola and Jones [12]. Instead of an exhaustive search over the whole MPR, only a local search is performed near the initial profile position extracted from the 3D face mesh candidate. To infer the optimal profile shape for the detected location, a weighted nearest-neighbor approach by Georgescu *et. al* [5] is used to find the closest profile prototype in the

training set. One can refer to [11, 12] and [5] for the technique details of rigid object detection and shape inference respectively.

3.3. Fetal Face Surface Deformation

After the localization of the 3D fetal face, we align the mean face mesh within the bounding box to get an initial face shape. Then we use a set of local boundary detectors to adjust the position of each mesh point. These boundary detectors are trained using PBT and steerable features around the mesh points, similarly to [8, 16]. Formally, a mesh point \hat{P} is selected to maximize the probability of being on the boundary B in volume V :

$$\hat{p} = \arg \max_{p \in Q_p} Pr(B|p, V), \quad (5)$$

where Q_p is a set of candidate points within the local search range around p . Unlike in [8, 16] where Q_p is along the normal direction of the boundary at p for a reduced computation, we search a surrounding region because of the sparsity of our mesh points. To achieve a smooth and natural face mesh, the adjusted shape is projected into the shape subspace to get a final detected face mesh.

4. 3D Fetal Face Carving

One goal of our system is to generate a pleasant rendering of the fetal face, which requires omitting all the obstructions in front of the face. Based on the detected fetal face mesh, we can achieve the goal by carving out all the voxels in front of the face. We develop a carving algorithm to perform this. The basic idea of the algorithm is simple. First we calculate all the surface points on the face mesh triangles, then along the reverse direction of the view, we set all the points on the rays from those surface points to have zero intensity. The carving algorithm is described in Algorithm 1.

5. Experiments

5.1. DataSet

Our experiments are performed on a dataset containing 1010 fetal face volumes, collected from 51 different fetuses at different pregnant periods from the 21st week to 40th week. The average size of the volumes is $157 \times 154 \times 104$ (mm) and the resolution is isotropically 1mm for all the volumes. All these volumes have been manually annotated with a face mesh. As mentioned in the introduction, this dataset is very challenging because of the large variation among the scanned data. For example, the fetuses at the 20th week of the pregnancy have quite different appearances with those at the 30th week. Different obstructions may exist in the volumes, e.g. extremities and cords, some

Algorithm 1 3D Face Carving

```
1:  $VolMask \leftarrow initialize\ false$ 
2:  $PointsInTriangles \leftarrow initialize\ empty$ 
3: for all mesh triangles  $T$  do
4:   collect points within  $T$  to  $PointsInTriangles$ 
5: end for
6: for all  $PointsInTriangles$   $p$  do
7:    $p_{front} = p$ 
8:   while  $p_{front}$  within volume do
9:      $voxel\ at\ p_{front} \leftarrow 0$ 
10:     $p_{front} - = faceNorm$ 
11:   end while
12:    $p_{back} = p$ 
13:   while  $p_{back}$  within volume do
14:      $VolMask[p_{back}] = 1$ 
15:      $p_{front} + = faceNorm$ 
16:   end while
17: end for
18:  $K \leftarrow calculate\ plane\ of\ face\ contour$ 
19: for all voxel indices  $q$  do
20:   if  $VolMask[q] == 0$  and  $voxel\ at\ q \neq 0$  and  $q$  in
     front of  $K$  then
21:      $voxel\ at\ q \leftarrow 0$ 
22:   end if
23: end for
```




Figure 6. The distributions of position, scale, orientation and week of the fetal faces in our dataset. Each point in (a) indicates the normalized x and y coordinates of a face center and each point in (b) indicates the width and height of a fetal face. In (c), each point represents the Euler angles of a fetal face rotated around x, y , and z axes respectively. (d) shows the distribution of the number of fetuses across different weeks.

of which may be directly on the faces. Furthermore, the position, orientation and size of fetal faces change a lot in the

data set. Figure 6 shows the distributions of position, scale, orientation and week of the fetal faces in our dataset.

In our experiments for 3D face mesh detection, out of the 1010 volumes 962 volumes are used for training and 48 volumes for testing. The test set consists of volumes of different fetuses and different pregnant weeks. The partition also guarantees that for each volume in the test set, there are no volumes in the training set that are from the same fetus and the same week. This setting makes sure that the performance on the test set can truly reflect the generalization of our models on totally unseen data. The 2D profile detector was trained on 500 annotated profile MPRs¹ from the ultrasound volumes, and tested on 48 profile MPRs from the volumes in the test set.

The features used for PBT training for face mesh translation were the 3D Haar features for its efficient computation using integral volumes. For orientation and scale training, we used the steerable features [16], which can not only capture the orientation and scale of the target object but also be computed very efficiently.

5.2. Detection Evaluation

In this section, we report the results for both the training and test set. The performance on the training set indicates the limit of our approach while that on the test set reflect how well the approach works on novel data. The performance metrics we used are based on the Euler distance of the detected results and the ground truth. For each test data we calculate the distances between corresponding points on the detected result and the ground truth, and take the average distance as an error measurement for that test data. We measure this error for all the test data and do statistics over all the error measurements.

To investigate how the detection models work, we evaluate the performances of the 2D profile detection and the 3D face mesh detection separately. Table 1 shows the performance of the 2D face profile detector on the 2D profile images cut from the volumes, from which we can see that the average distance of the detected profile points with the ground truth is about 3.4 mm and for the test set this number is slightly higher (3.7 mm). The profile detector performs very fast, only 0.3 seconds per image of 512×512 .

Table 2 shows the performance of the 3D face mesh detection on training set and test set respectively. Through combining both 3D information and 2D profile information, the average error of the fetal face detection on test set is improved by more than one pixel. The big improvement of error standard deviation from 10.5 to 4.2 shows the combined approach is much more robust than the pure 3D mesh detection. Figure 8 shows some results of 3D face mesh detection with different errors.

¹This is the number of volumes we had when training the profile detector

2D profile	Mean (mm)	Median (mm)	Std (mm)	Run Time (seconds)
Training Error	3.4	2.9	2.0	0.3
Test Error	3.7	3.5	1.4	0.3

Table 1. Performance on 2D face profile detection over the training set and test set

3D Mesh	Mean (mm)	Median (mm)	Std (mm)	Run Time (seconds)
Training Error	4.3	3.6	3.5	0.8
Test Error	8.8	6.2	10.5	0.8
Test Error After 2D refinement	7.6	6.6	4.2	1.0

Table 2. Performance on 3D fetal face surface detection over the training set and test set. A sign test shows that the performance improvement by using 2D refinement is statistically significant with the P-value equal to 0.03.




Figure 7. Examples of 2D profile detection. The first column shows the input cutting image from ultrasound volume, and the second and the third columns show our detected results and human annotation respectively.

5.3. Face Carving Based on the Detected Mesh

Figure 8 shows examples of 3D carving based on the 3D detection results with different detection errors. From the second, third and fourth columns in the figure, we can see that although the 3D detection obtains a correct view for the fetal face, we do not obtain a satisfying rendering because of the clouds and occlusions in front of them. Through carving, the obstructions in front of the face are successfully removed and a much better 3D view is rendered as shown in the figure. The carving example at the bottom row of this figure is based on the detected face mesh with a 6.22mm error, which is still visually acceptable.

6. Conclusions and Future Work

In this paper, we propose an efficient and robust learning based approach for fetal face detection from 3D ultrasound data. Our approach combines both 3D face surface detection and 2D face profile detection to find the optimal acquisition plane for face rendering. We use constrained marginal space learning for 3D face surface detection and a boosting

approach for 2D profile detection. To enhance the 3D fetal face rendering, we propose a carving algorithm to remove all obstructions in front of the face based on the detection results. Our experiments show excellent detection accuracy and fast speed of the system on a large fetus data set. In the future, we will further improve the detection performance through exploiting more sophisticated features for 3D face box detection and geometric constraints among landmarks for boundary adjustment. We will also explore mesh based transparency adjustment for better and smoother face rendering.

References

- [1] C. Beumier and M. Acheroy. Face verification from 3D and grey level clues. *Pattern Recognition Letters*, 12:1321–1329, 2001.
- [2] G. Carneiro, F. Amat, B. Georgescu, S. Good, and D. Comaniciu. Semantic-based indexing of fetal anatomies from 3-D ultrasound data using global/semi-local context and sequential sampling. In *CVPR*, 2008.
- [3] J. Cartoux, J. Lapreste, and M. Richetin. Face authentication or recognition by profile extraction from range images. In *IEEE Computer Society Workshop on Interpretation of 3D Scenes*, pages 194–199, 1989.
- [4] A. Colombo, C. Cusano, and R. Schettini. 3D face detection using curvature analysis. *Pattern Recognition*, 2006.
- [5] B. Georgescuand, X. Zhou, D. Comaniciu, and A. Gupta. Database-guided segmentation of anatomical structures with complex appearance. In *CVPR*, pages 429–436, 2005.
- [6] J. Huang, B. Heisele, and V. Blanz, editors. *Component-based face recognition with 3D morphable models*, chapter Audio- and Video-Based Biometric Person Authentication. Springer Berlin, 2003.
- [7] C. F. F. Karney. Quaternions in molecular modeling. *Journal of Molecular Graphics and Modelling*, 25:595–604, 2007.
- [8] H. Ling, S. K. Zhou, Y. Zheng, B. Georgescu, M. Suehling, and D. Comaniciu. Hierarchical, learning-based automatic liver segmentation. In *CVPR*, 2008.
- [9] A. Moreno, A. Sanchez, J. Velez, and F. Diaz. Face recognition using 3D surface extracted descriptors. In *Proceedings of the Irish Machine Vision and Image Processing 2003, 2004*, (ISBN 1-85923-177-2).
- [10] G. Pan, Y. Wu, and Z. Wu. Investigating profile extracted from range data for 3D face recognition. In *IEEE International Conference on Systems, Man and Cybernetics.*, pages 1396–1399, 2003.
- [11] Z. Tu. Probabilistic boosting-tree: Learning discriminative models for classification, recognition, and clustering. In *ICCV*, pages 1589–1596, 2005.
- [12] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. In *CVPR*, 2001.
- [13] Y. Wang, C. Chua, and Y. Ho. Facial feature detection and face recognition from 2D and 3D images. *Pattern Recognition Letters*, pages 1191–1202, 2002.
- [14] M. Yang. Recent advances in face detection. *IEEE ICPR 2004 Tutorial*, 2004.




Figure 8. Examples of 3D fetal face detection and carving. From left to right: original position of loaded volumes, detected face meshes, views before carving based on the detected meshes, views with an angle, views after carving based on the detected meshes, views with an angle to see the carving effects. The 3D detection errors from top to bottom: 2mm, 3mm, 4.25mm, 5mm, 6.22mm.

[15] M. Yang, D. Kriegman, and N. Ahuja. Detecting faces in images: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 24(1):34–58, 2002.

[16] Y. Zheng, A. Barbu, B. Georgescu, M. Scheuring, and D. Comaniciu. Fast automatic heart chamber segmentation from 3D CT data using marginal space learning and steerable features. In *ICCV*, 2007.

[17] Y. Zheng, B. Georgescu, H. Ling, S. K. Zhou, M. Scheuer-

ing, and D. Comaniciu. Constrained marginal space learning for efficient 3D anatomical structure detection in medical images. In *CVPR*, 2009.